Course Syllabus for ELEC 677: Neural Methods for Image Synthesis
Spring 2022
Instructor: Guha Balakrishnan (guha@rice.edu)

Overview

This course will cover state-of-the-art research in image synthesis methods using neural networks. Lectures will mainly revolve around discussing important publications in the field.

Recommended background

Basic knowledge of deep learning and computer vision. Experience with Python is also needed for the final project.

Textbooks

None, though the free Deep Learning Book and Computer Vision: Algorithms and Applications are useful resources.

Assignments & Grading

Class Participation (30%):
Students should submit 2 questions they would like answered/discussed related to each lecture's main paper by 5PM the day before the lecture. Questions should be posted to this Google doc, with your name. The presenter for that lecture can use these for discussion and planning purposes. In addition to these submitted comments, students are expected to actively participate during lectures.

Student Presentation (35%):
Each student is responsible for leading one lecture during the semester. This presentation will cover the full details of a single paper as well as any related work needed to give context to the paper. The presenter is also responsible for fostering discussion during the presentation, in part by using the submitted questions from classmates (see class participation above).

Final Project (35%):
Students will work in teams of 1-3 to complete a final coding-based project of their own design. Each project will do one of the following:

1. Design a novel tweak or modification to an existing algorithm, code it, and show results/performance.
2. Apply an existing algorithm to a novel application not shown in the original paper.
3. Thoroughly analyze an existing algorithm empirically and/or theoretically, highlighting pros and cons.

4. Design + code a new algorithm for an established (or novel) synthesis task (hardest project, but possible with some creativity!)

Topics Covered

Lectures will cover the following topics (in roughly the same order):

1. Course outline/intro
2. Basics of CNNs, convolutions and image operations
3. Style Transfer and Texture Synthesis
4. Perceptual Losses
5. Deep Image Priors
6. Intro to generative models and GANs
7. StyleGAN2
8. Image-to-image translation
9. Image editing and traversals with GANs
10. Unpaired image translation
11. Towards 3D-aware GANs
12. Training GANs from a single image
13. Inferring motion from a single image
14. Video generation
15. Variational Autoencoders
16. Intro to graphics basics: light transport, rendering, intrinsic images
17. Neural scene representations
18. Novel View Synthesis
19. Compositional neural fields
20. Inverse rendering
21. Relighting
22. Face Animation
23. Ethics